A Benchmark and Comparison of Active Learning for Logistic Regression

Yazhou Yang Y.YANG-4@TUDELFT.NL
Marco Loog M.LOOG@TUDELFT.NL
Pattern Recognition Laboratory, Delft University of Technology, Mekelweg 4, 2628CD Delft, The Netherlands

arXiv:1611.08618v1 [stat.ML] 25 Nov 2016

Abstract

Various active learning methods based on logistic regression have been proposed. In this paper, we investigate seven state-of-the-art strategies, present an extensive benchmark, and provide a better understanding of their underlying characteristics. Experiments are carried out both on 3 synthetic datasets and 43 real-world datasets, providing insights into the behaviour of these active learning methods with respect to classification accuracy and their computational cost.

1. Introduction

In practice, acquiring a large amount of data is easy while labeling these data is often rather difficult, time-consuming or expensive, due to the involvement of human experts (Settles, 2010). For instance, collecting millions of images from Google is not that difficult, while obtaining the categories of each image may need a lot of manpower and other resources. Active learning addresses this challenge by selecting a most valuable subset from the whole data set for human annotation. The key idea behind is that active learning has the abilities and freedom to automatically choose the most valuable instances according to some criterion. Many research studies have demonstrated that active learning can be effective in maintaining good performance while reducing the overall labeling effort over a diverse range of applications, such as text categorization (Tong & Koller, 2002; Cai & He, 2012), medical image classification (Hoi et al., 2006), image retrieval (Hoi et al., 2008) and natural language processing (Tang et al., 2002).

Choosing an appropriate criterion which measures the usefulness of unlabeled instances is of vital importance for active learning. Active learning tends to select those instances which can provide as much information as possible for the learning problem. Perhaps, the simplest method to choose the next instance from the unlabeled pool is random sampling. Even though that random sampling strategy is expected to be less efficient than actual active learning algorithms, at times, it can perform very well in comparison with the latter (Guyon et al., 2011). Most commonly used criteria in active learning include query-by-committee (Seung et al., 1992), uncertainty sampling (Tong & Koller, 2002), expected error minimization (Roy & McCallum, 2001; Guo & Greiner, 2007; Holub et al., 2008), and variance reduction (Zhang & Oles, 2000; Yu et al., 2006; Schein & Ungar, 2007), model change (Settles et al., 2008; Freytag et al., 2014; Cai et al., 2014) and "min-max" view active learning (Hoi et al., 2008; Huang et al., 2014). They are derived from diverse heuristics and based on various types of classifiers. Some of them are specifically designed for one certain classifier, *e.g.*, the simple margin criterion for support vector machines (Tong & Koller, 2002), while others can be adapted to various classifiers, *e.g.*, expected error reduction works with logistic regression and naive Bayes (Roy & McCallum, 2001).

Since the base classifier is a key component for an active learning algorithm, it seems fair to compare active learning methods using the same base classifier. In this paper, we mainly focus on several active learning approaches derived for logistic regression model. Logistic regression is chosen because (1) logistic regression the most widely applied classifier, certainly outside the areas of machine learning and pattern recognition. (2) For logistic regression, an estimate of the posterior class probability is readily available, something that is of use in active learning. It is, in part, for this reason that most active learners employ this methods as base classifier and a more diverse set of active learners have been developed. (3) there are already a variety of active learning algorithms which have been developed specifically for logistic regression (Kanamori & Shimodaira, 2003; Hoi et al., 2006; Guo & Greiner, 2007; Schein & Ungar, 2007; Kanamori, 2007; Cuong et al., 2014; Liu et al., 2015).

Hence, in this work, we evaluate active learning methods which can be used in combination with logistic regression. As we see it, the various methods can be considered to be based on four different heuristics: uncertainty sample, error reduction, variance reduction and min-max view active

learning, all of which will be explained in further detail in Section 2. Seven different active learning algorithms within the above heuristics are compared. We aim to better understand their underlying characteristics and how these methods work on classification tasks with different levels of difficulty and why. We should mention the work by Schein and Ungar (Schein & Ungar, 2007) here, who already did an evaluation of active learning methods using logistic regression. In this paper, however, we compare some new methods, which appeared only recently (Guo & Greiner, 2007; Guo & Schuurmans, 2008; Huang et al., 2014), and we generally provide an evaluation much more extensive than their paper does. We also investigate how active learning algorithms generally perform in comparisons with random sampling and point out the underlying relationships among the compared methods. The computational cost of each method is also evaluated.

In this paper, we focus on the setting of pool-based active learning that assumes a large pool of unlabeled data along with a small set of labeled data that is already available (Settles, 2010). We consider the myopic setting which selects an unlabeled instance at a time. Batch mode active learning, which selects a batch of examples simultaneously, is not considered in this work and we refer for further background and some typical approaches to (Hoi et al., 2006; Guo & Schuurmans, 2008; Chakraborty et al., 2010; 2014).

The main contributions of this work can be summarized as follows:

1. A review of the state-of-the-art in active learning for logistic regression is presented in which links and relationships between methods are explicated.
2. A preference map is proposed that reveals characteristics of the selection location in 2D problems.
3. Extensive experiments on 43 real-world data sets and 3 artificial sets are carried out.
4. Insights are provide into the behaviour of the classification accuracy and computational cost.

1.1. Outline

The remainder of the paper is organized as follows. Section 2 firstly describes the general procedure of active learning, and then reviews the various approaches to active learning built on logistic regression. At the same time it sketches the relationships between the different methods. Extensive experiments on synthetic and real-world data sets are given in Section 3. The experimental setup is described and the outcomes are reported on. More importantly, it provides an extensive discussion of the findings and aims to critically evaluate these. Section 4 concludes the work.

2. Active Learning Strategies and Methods

For myopic active learning in the pool-based scenario, we often assume that a small set of labeled instances with a large pool of unlabeled samples are available. Let $\mathcal{L} = \{(x_i, y_i)\}_{i=1}^{l}$ represent the training data set that consists of l labeled instances and $\mathcal{U}$ be the pool of unlabeled instances $\{x_i\}_{i=l+1}^{n}$. Each $x_i \in \mathbb{R}^d$ is a d dimensional feature vector, and $y_i \in C$ is the class label of x_i. In this work we restrict ourselves to binary classification, which does not pose any essential limitation. For this reason, C is simply taken to be the set $\{+1, -1\}$. We denote $P_{\mathcal{L}}(y|x)$ be the conditional probability of y given x according to a logistic classifier trained on $\mathcal{L}$. The active learner will select a best instance x_i from the unlabeled pool based on its measure of utility, ask for the corresponding label y_i from the human and extend the training set with the new labeled sample $\mathcal{L} = \mathcal{L} \cup (x_i, y_i)$. The whole procedure is repeated many times until some stopping criterion is satisfied.

The remaining part of this section presents seven active learning algorithms built on logistic regression. These seven algorithms relate to four different heuristics, one per subsection, that drive the design of the active learner. The links between some active learners are also presented.

2.1. Maximum Entropy

One widely used approach to active learning is uncertainty sampling (Lewis & Gale, 1994; Settles, 2010). It selects the instances for which the current classifier is least certain. Querying these least certain, or equivalently, easily misclassified instances can help the model clarify the position of decision boundaries. Intuitively, the distances between unlabeled instances and the decision boundary can be a measure of the uncertainty. Tong *et al.* (Tong & Koller, 2002) proposed a simple margin approach which indeed queries the instance closest to the decision boundary.

Entropy (Shannon, 1948) is a different and more widely used general measure of uncertainty. Entropy-based approaches query the instances with *maximum entropy*:

$$\arg\max_x -\sum_{y \in C} P_{\mathcal{L}}(y|x) \log P_{\mathcal{L}}(y|x) \qquad (1)$$

This method is called ENTROPY for short. It calculates the entropy of each $x \in \mathcal{U}$ and selects the instance which has maximum entropy. It can be used with any classifier that produces probabilistic output. For binary classification, ENTROPY is equivalent to the simple margin approach (Settles, 2010).

One shortcoming of such uncertainty sampling based approaches lies in that they are easy to pick up outliers since outliers are likely to acquire high uncertainty values (Tang et al., 2002; Zhu et al., 2010). It takes a risk being stuck

in a situation where active learner keeps selecting instances which do not improve the current classifier at all (Huang et al., 2014).

2.2. Expected Error Reduction

Expected error reduction approaches are another type of popular active learning methods (Roy & McCallum, 2001; Guo & Greiner, 2007; Holub et al., 2008; Guo & Schuurmans, 2008). These approaches attempt to measure how much the generalization error is likely to be reduced when adding one new instance into the labeled data set. As one does not have direct access to the future test data, however, Roy and McCallum (Roy & McCallum, 2001) proposed to estimate the expected error rate over the unlabeled examples under the assumption that the unlabeled data set is representative of the test distribution. In other words, the unlabeled pool can be viewed as a validation set. Roy and McCallum proposed to estimate the expected error using expected log-loss or 0-1 loss. For the former, which we consider in our work, the following objective is considered:

$$\arg\min_{x\in\mathcal{U}} \sum_{y\in C} P_{\mathcal{L}}(y|x) \left(- \sum_{x_i\in\mathcal{U}} \sum_{y_i\in C} P_{\mathcal{L}^+}(y_i|x_i) \log P_{\mathcal{L}^+}(y_i|x_i) \right) \tag{2}$$

Here, $\mathcal{L}^+ = \mathcal{L} \cup (x, y)$ indicates that the selected instance x is labeled y and added to the labeled data set $\mathcal{L}$. We refer to this method as EER in this paper. The first term $P_{\mathcal{L}}(y|x)$ is the posterior probability of y given x trained on the labeled data set $\mathcal{L}$. The second term can be viewed as the sum of entropy or (equivalently) expected log-likelihood over the unlabeled data $\mathcal{U}$.

However, since the labeled data $\mathcal{L}$, which is typically of small size, might result in a bad classifier, $P_{\mathcal{L}}(y|x)$ may be not helpful (Guo & Greiner, 2007). To avoid problems with such misspecifications, Guo and Greiner (Guo & Greiner, 2007) proposed an optimistic or, equivalently, *maximum* error reduction approach (called MAXER in this paper), which estimates the best-case error reduction, without considering $P_{\mathcal{L}}(y|x)$. MAXER considers the following objective function instead:

$$\arg\min_{x\in\mathcal{U}} \min_{y\in C} \left(- \sum_{x_i\in\mathcal{U}} \sum_{y_i\in C} P_{\mathcal{L}^+}(y_i|x_i) \log P_{\mathcal{L}^+}(y_i|x_i) \right) \tag{3}$$

Note that the error reduction approaches above only take the unlabeled instances into consideration and fail to include labeled examples. Therefore, in order to obtain better generalization performance, it has been suggested to compute the loss over the training set $\mathcal{L}$, while at the same time estimate the loss over the unlabeled set $\mathcal{U}$. This idea is proposed in (Grandvalet & Bengio, 2004) for semi-supervised learning, while Guo (Guo & Schuurmans, 2008) extends it to batch mode active learning. As we focus on myopic active learning, we adopt the following related criterion:

$$\arg\min_{x\in\mathcal{U}} \min_{y\in C} - \left(\sum_{x_j\in\mathcal{L}^+} \log P_{\mathcal{L}^+}(y_j|x_j) + \alpha \sum_{x_i\in\mathcal{U}\setminus x} \sum_{y_i\in C} P_{\mathcal{L}^+}(y_i|x_i) \log P_{\mathcal{L}^+}(y_i|x_i) \right) \tag{4}$$

where α is a trade-off parameter used to adjust the importance of loss over labeled and unlabeled data. We term this combined approach CEER in this paper. This approach can of course be adopted directly to the worst-case or average-case criterion as well.

One general disadvantage of error reduction approaches may be the high computational cost (Settles, 2010). Each time a new queried instance x with its label y is added to the training data set, we need to retrain the classifier to get the new posterior probabilities $P_{\mathcal{L}^+}(y_i|x_i)$. This retraining step may amount to great computational efforts.

2.3. Variance Reduction

Optimal experimental design (OED) is a next approach to active learning that we consider. In general, it attempts to minimize some statistical criteria with the aim of saving cost of an experimental design (Atkinson et al., 2007; Yu et al., 2006). For example, A-optimality is one of the commonly used measure and it seeks to minimize the trace of the inverse of the information matrix (Atkinson et al., 2007). Equivalently, A-optimality can also be seen as reducing the average variance of estimated parameter and is wildly used in active learning (Zhang & Oles, 2000; Hoi et al., 2006; Schein & Ungar, 2007). Experimental design has a close relationship with maximum likelihood estimation (MLE). The Fisher information matrix is the second derivative of likelihood function with regards to model parameters. It is defined as $\mathcal{I}(\theta) = -E[\frac{\partial^2}{\partial\theta^2} \log f(X|\theta)]$, where θ contains all model parameters and $f(X|\theta)$ is the likelihood function. The observed Fisher information matrix is $\mathcal{I}(\hat{\theta}_{MLE})$: the information matrix evaluated at the maximum likelihood estimate.

In the binary setting, for regularized logistic regression, we consider minimizing $\min \sum_{i\in\mathcal{L}} \log(1+\exp^{-y_i w^T x_i}) + \frac{1}{2\lambda}||w||^2$, where w is the parameter of logistic regression model, λ is the regularization parameter and $y_i \in \{+1, -1\}$. In this case, the Fisher information matrix over the unlabeled pool $\mathcal{U}$ is defined as $\mathcal{I}_{\mathcal{U}}(w) = \frac{1}{|\mathcal{U}|} \sum_{i\in\mathcal{U}} \sigma_i(1-\sigma_i) x_i x_i^T + \lambda I_d$ where $\sigma_i = \sigma(w^T x_i) = 1/(1+\exp(-w^T x_i))$ is the posterior probability of $P(y = 1|x_i)$, and I_d is the identity matrix of size $d \times d$. Zhang

et al. (Zhang & Oles, 2000) utilize A-optimal design and minimize the Fisher information ratio between $\mathcal{I}_{\mathcal{U}}(\hat{w})$ and $\mathcal{I}_x(\hat{w})$:

$$\arg\min_{x\in\mathcal{U}} \mathrm{tr}(\mathcal{I}_x(\hat{w})^{-1}\mathcal{I}_{\mathcal{U}}(\hat{w})) \tag{5}$$

where $\mathcal{I}_x(\hat{w}) = \sigma_i(1-\sigma_i)x_i x_i^T + \lambda I_d$ and $\hat{w}$ is the maximum likelihood estimator. The entity $\mathcal{I}_{\mathcal{U}}(\hat{w})$ can be interpreted as the model's output variance with respect to unlabeled data $\mathcal{U}$, and $\mathcal{I}_x(\hat{w})^{-1}\mathcal{I}_{\mathcal{U}}(\hat{w})$ can be viewed as the future output variance once x has been labeled. The criterion suggested selects unlabeled examples that minimize this ratio or, equivalently, reduce the future variance. We call this approach Fisher information variance reduction (FIVR) in this paper. Hoi *et al.* (Hoi et al., 2006) exploited the same idea as in (Zhang & Oles, 2000) and extended it to the batch mode setting. When the batch size is set to one, Hoi's method is basically identical to FIVR apart from some approximations introduced to be able to deal with the batch setting.

Schein and Ungar (Schein & Ungar, 2007) proposed a similar A-optimal active learning method based on logistic regression. In doing so, one can define the Fisher information matrix over the training data set $\mathcal{L}$ as $F = \frac{1}{l}\sum_{i\in\mathcal{L}} \sigma_i(1-\sigma_i)x_i x_i^T + \lambda I_d$. Schein proposed to minimize the variance of the estimated distribution of the model's output $\sigma(\hat{w}^T x_i)$ and considers a Taylor expansion around $\sigma(\hat{w}^T x_i)$:

$$\sigma(\hat{w}^T x_i) \simeq \sigma(w^T x_i) + c_i^T(\hat{w} - w)$$

where $c_i = \sigma_i(1-\sigma_i)x_i$ is the gradient vector of σ_i. Based in this, Schein notes that

$$\mathrm{Var}(\sigma(\hat{w}^T x_i)) \simeq \mathrm{Var}(c_i^T(\hat{w} - w)) \simeq c_i^T F^{-1} c_i$$

Note that $\mathrm{Var}(\sigma(\hat{w}^T x_i))$ is equivalent to $\mathrm{Var}(\sigma(-\hat{w}^T x_i))$. The variance over all the unlabeled instance can be formulated as follows:

$$\begin{aligned} g(\mathcal{L},\mathcal{U}) &= \sum_{x_i\in\mathcal{U}} \sum_{y\in\{+1,-1\}} \mathrm{Var}(\sigma(y\hat{w}^T x_i)) \\ &\simeq 2\sum_{x_i\in\mathcal{U}} tr\{c_i^T F^{-1} c_i\} \end{aligned} \tag{6}$$

The benefit of a newly selected instance, can then be measured in terms of the expected variance reduction:

$$\arg\min_{x\in\mathcal{U}} \sum_{y\in C} P_{\mathcal{L}}(y|x)\, g(\mathcal{L}\cup(x,y),\mathcal{U}) \tag{7}$$

We refer to this method as EVR in this paper. EVR represents the potential variance changes weighted by current estimated model $P_{\mathcal{L}}(y|x)$. EVR can also be extended to log-loss based EVR (Schein & Ungar, 2007). We ignore this algorithm since we observe that it performs poorer than EVR in our experiments.

EVR also suffers from the problem of large computational costs since it goes over all the pool and re-estimates $\hat{w}$ and Fisher information matrix each time. The calculation of the inverse of matrix even makes it slower than expected error reduction approaches.

2.4. Minimum Loss Increase

The next heuristic we consider is minimum loss increase, which directly bases its criterion on already labeled samples. Related to this class, is the method originally proposed by Hoi (Hoi et al., 2008) that minimizes the gain of the objective function, which is also referred to as a "min-max" (or "worst-case") view on active learning. We here look at this work in a more general formulation and demonstrate its relationship with the expected error reduction framework.

Let us consider an unconstrained optimization problem using an L2-loss regularized linear classifier and a loss function $V(w; x_i, y_i)$:

$$\min_w \qquad g(w) = \frac{1}{2\lambda}||w||^2 + \sum_{i=1}^{l} V(w; x_i, y_i)$$

where $y \in \{+1, -1\}$. Many loss functions can be adopted for linear classification. For example, hinge loss, $V(w; x_i, y_i) = \max(0, 1 - y_i w^T x_i)$, results in linear SVM and squared loss, $V(w; x_i, y_i) = (y_i - w^T x_i)^2$, leads to ridge regression. We, of course, will consider the logistic loss in the experimental section: $V(w; x_i, y_i) = \log(1 + \exp^{-y_i w^T x_i})$, which results in L2-regularized logistic regression. Note that minimizing the logistic loss over the training data is equivalent to maximizing the conditional likelihood.

Now, in order to identify the most valuable instances for labeling, we could select the example that, once labeled, results in the minimum gain in terms of the value of objective function. That is, we consider

$$\arg\min_{x\in\mathcal{U}} \max_{y\in C} g_{\mathcal{L}^+}(w) - g_{\mathcal{L}}(w) \tag{8}$$

where $g_{\mathcal{L}}(w)$ denotes the score of objective function over the training data $\mathcal{L}$. The second term, $g_{\mathcal{L}}(w)$, is independent of the next queried instance and we can rewrite Equation 8 as follows:

$$\arg\min_{x\in\mathcal{U}} \max_{y\in C} \min_w \frac{1}{2\lambda}||w||^2 + \sum_{x_i\in\mathcal{L}^+} V(w; x_i, y_i) \tag{9}$$

This method can be interpreted as directly minimizing the worst-case value of the objective function when labeling a new instance.

Considering kernel versions instead of linear classifiers in the above, would entail the earlier mentioned min-max active learning (Hoi et al., 2008; Huang et al., 2014), which

use the hinge loss and square loss, respectively. Reference (Hoi et al., 2008) originally presented the min-max view and extended it to batch mode active learning. It also shows that the approach in some sense approximates margin methods. Huang (Huang et al., 2014) proposed so-called QUIRE which extends this framework to consider all the unlabeled data to select both informative and representative instances in Equation 10.

$$\arg\min_{x\in\mathcal{U}} \min_{y_u\in C^{\{n_u-1\}}} \max_{y\in C} \min_{w} \frac{1}{2\lambda}||w||^2 + \sum_{x_i\in\mathcal{L}\cup\mathcal{U}} V(w;x_i,y_i) \tag{10}$$

where y_u indicates the labels of remaining unlabeled pool $\mathcal{U}\backslash x$ since x is already selected and added to $\mathcal{L}^+$, and n_u is the number of samples of U. We need to point out, however, that this unlabeled part $\mathcal{U}\backslash x$ is of no use since QUIRE relaxes the constraints of y_u. Optimizing this unconstrained y_u can guarantee that the remaining unlabeled data $\mathcal{U}\backslash x$ is useless, which can also be observed from equation (9) and (10) in QUIRE paper (Huang et al., 2014). Therefore, QUIRE also fits our general framework.

As we consider the logistic loss for the above framework, we will refer to this method with MLI since the method aims to minimize the increase of the value of an objective function. The objective, we consider is as follows:

$$\begin{aligned}
&\arg\min_{x\in\mathcal{U}} \max_{y\in C} \min_{w} \frac{1}{2\lambda}||w||^2 + \sum_{x_i\in\mathcal{L}^+} V(w;x_i,y_i) \\
&= \arg\min_{x\in\mathcal{U}} \max_{y\in C} \min_{w} \frac{1}{2\lambda}||w||^2 + \sum_{x_i\in\mathcal{L}^+} \log(1+\exp^{-y_i w^T x_i}) \\
&= \arg\min_{x\in\mathcal{U}} \max_{y\in C} \frac{1}{2\lambda}||\hat{w}||^2 - \sum_{x_i\in\mathcal{L}^+} \log P_{\mathcal{L}^+}(y_i|x_i)
\end{aligned} \tag{11}$$

where $\hat{w}$ is the estimated parameter of the L2-regularized logistic regression model trained on the labeled data $\mathcal{L}^+ = \mathcal{L}\cup(x,y)$.

2.5. Links and Connections

The above active learning methods are all constructed using logistic regression, but some further relationships can be identified, as shown in Table 1. We summarize seven active learning algorithms and analyse their links and connections. "data" in Table 1 indicates which kind of data is used to evaluate the criterion of an instance. For example, ENTROPY and FIVR measure their criteria only on the individual instance itself while EER calculates its criterion on all unlabeled data. "scenarios" represents which kind of scenario is considered to deal with missing true label when retraining classifier, *e.g.*, average-case scenario for EER and worst-case scenario for MLI.

We see that EER, MAXER, CEER, MLI and EVR are rather similar in some respects. First, they all belong to retraining-based active learning, which measures the utility of instance by repeatedly labeling it (any possible category) and retraining the model. Since the true label is not available before asking for human annotation, these methods turns to consider the average-case, worse-case or best-case scenarios. For example, both EER and EVR adopt the average-case criterion, or equivalently, the expectation value. Similarly, MAXER and CEER choose the best-case value while MLI takes the worst-case situation into consideration. Second, their measure of usefulness of unlabeled samples are derived on different kind of data. For example, EER, EVR and MAXER evaluate their corresponding criteria only on the unlabeled set while MLI does the evaluation on labeled set. CEER considers all the available data. Third, they evaluate the utility of samples by different means. EER and MAXER compute the sum of entropy as a measure of error while EVR cares about the variance. MLI bases its judgement on the value of objective function. CEER can be seen as a combination of EER and MLI, which measures the entropy and log-likelihood simultaneously with, in addition, a small difference of using the best-case criterion and not using any regularization.

ENTROPY and FIVR are quite different from the above five active learning algorithms since they both do not need retrain the classifier multiple times during one selection. ENTROPY favours the instances with maximum entropy without retraining the model, while EER considers the sum of entropy over all unlabeled data after retraining the model. FIVR, which aims to reduce the Fisher information variance, is a little different from the EVR due to no need of retraining. Similarities and differences in performance between these active learning algorithms will be discussed in the next section, Section 3, which covers the experimental setup as well.

3. Experiments

In this section, we firstly describe the experimental settings. Then we analyse the experimental results on synthetic datasets and real-world datasets, respectively, followed by the computational cost analysis.

3.1. Experimental Setting

We mainly present the detailed information of 3 synthetic data sets and 43 real-world data sets in the following subsection, along with an illustration of the evaluation design.

3.1.1. Synthetic Data Sets

Three binary synthetic data sets are constructed to intuitively demonstrate the different behaviours of active learning algorithms. The first data set *Synth1* is a standard *2D*

Table 1. Links and connections among different active learning algorithms. "data" indicates which kind of data is used to evaluate the criterion of an instance; "scenarios" represents which kind of scenario is considered to deal with missing true label when retraining classifier.

heuristics	methods	data	scenarios
uncertainty sampling	ENTROPY	individual instance	-
error reduction	EER	unlabeled set	average-case
	MAXER	unlabeled set	best-case
	CEER	labeled+unlabeled set	best-case
min-max	EER	labeled set	worst-case
variance reduction	EVR	unlabeled set	average-case
	FIVR	individual instance	-

binary problem which is shown in Fig 1a. Positive and negative classes are generated according to two multivariate normal distributions centered at $[1, 1]^T$ and $[-1, -1]^T$, respectively. We want to explore which active learning method works well on this unambiguously specified problem. The second data set, *Synth2*, displayed in Fig 1d, is generated according to the description in (Huang et al., 2014). We can observe that Synth2 has a clear cluster structure. On this kind of data, uncertainty sampling has substantial problems since it only considers the most *uncertain* instance which comes closest to the decision boundary. Initially, the decision boundary estimated from the limited number of labeled data may be far away from the true one and therefore uncertainty sampling may select less informative instances due to a poorly estimated posterior probability. Of course, this is precisely what this data set was designed for and set out to illustrate. This data set prefers some kind of active learning methods which can consider the so-called representativeness along with the informativeness at the same time (Huang et al., 2014). Representative instances are those that lead more to exploration and not that much to exploitation. The latter is what uncertainty sampling typically aims for. The third synthetic data set, named *Synth3*, is also a 2D classification problem which is shown in Fig 1g. Synth3 is constructed to have a shape which looks like a tilted ⊔. Each part is generated from two multivariate normal distributions with small overlap. Compared with Synth1, Synth3 is a more challenging data set since it has relatively complex structure and may mislead some active learning methods. We are curious whether active learning can outperform random sampling on this kind of data. We investigate how active learning approaches work in the above three synthetic data sets and whether they perform better than random sampling.

In addition, we also explore the effect of the class overlapping and feature dimensionality on the performance of active learning algorithms. We consider a further experiment on Synth1 dataset by changing the overlap rates of two classes. This is carried out by decreasing the distance between the mean of two Gaussian distributions. We also adjust the dimensionality of Synth1 dataset from 2 to 20 and 200 dimensions to investigate the influence of feature dimensionality. The details can be found in Section 3.2.

3.1.2. Real-world Data Sets

As real-world benchmarks, we used various UCI data sets (Lichman, 2013), the MNIST handwritten digit dataset (LeCun et al., 1998) and the 20 Newsgroups dataset (Lang, 1995). Table 2 lists the data sets used in our study together with some basic information. All the datasets are pre-processed to become binary classification problems.

There are a total of 43 benchmark datasets included in this comparison. Most data sets are pre-processed to have zero mean and unity standard deviation according to (Fernández-Delgado et al., 2014). Some data sets are linearly scaled to $[-1, 1]$ or $[0, 1]$ according to (Chang & Lin, 2011) [1]. These datasets have various sample sizes and diverse feature dimensionalities. Some of them can be quite easily solved while others are quite difficult classification problems. The Letter Recognition Data Set from UCI, which consists of 20,000 examples of 26 uppercase letters in various fonts and distortions, is also used as a test bed in (Frey & Slate, 1991). As in this last work, 16 attributes are extracted from the letters as the feature and we consider the following six classification tasks between pairs of letters: D vs. P, E vs. F, I vs. J , M vs. N, V vs. Y, and U vs. V. These pairs of letters are selected since they can have a fairly similar appearance and distinguishing them can be challenging.

The MNIST [2] contains 60,000 training examples and 10,000 test examples which have been pre-processed to the same size of 28×28 pixels. The pairs 3 vs. 5, 5 vs. 8, and 7 vs. 9 constitute three difficult classification tasks and are used as the binary sets in our benchmark. For each of the three pairs, we randomly subsample 1500 instances

[1] https://www.csie.ntu.edu.tw/ cjlin/libsvmtools/datasets/
[2] http://yann.lecun.com/exdb/mnist/

from the original data set for computational reasons. Each pixel value is extracted as a feature, resulting in a 784-dimensional feature.

The 20 Newsgroups is a common benchmark used for document classification [3]. We use one version of this dataset which consists of 18,846 instances of 20 different news topics. Following the work of (Zhu et al., 2003), we also evaluate three binary tasks from this data set: sport.baseball vs. sport.hockey, pc.hardware vs. mac.hardware, and talk.religion.misc vs. alt.atheism. All the documents have been pre-processed into 26,241 dimensional tf.idf vectors to which we initially apply PCA to reduce the dimensionality to 500 for computational reasons.

3.1.3. Evaluation design

In the evaluation, each data set is randomly divided into training and test data sets of equal size. Following previous work (McCallumzy & Nigamy, 1998; Zhu et al., 2003; Tong & Koller, 2002; Guo, 2010; Gu et al., 2014; Li et al., 2012), we consider a difficult case of active learning, where only two labeled instances are provided as the initial labeled set, one from each class. We repeat each experiment 20 times on each real-world dataset. As to the synthetic data sets, we repeat the experiments 1000 times and every time we randomly regenerate the whole data set. The average performance of each active learning method on each data set is reported. In all the experiments, regularized logistic regression included in LIBLINEAR package (Fan et al., 2008) is used as the classifier. We set the regularization parameter λ to 0.01. The trade-off parameter present in the active learners considered, α for CEER, is set to 1.

For performance comparison, classification accuracy (or basically equivalently, the error rate) is the defacto standard evaluation criterion: the higher the accuracy, the better the algorithm. In active learning, however, performance varies depending on the number of labeled samples that one is allowed to take and we cannot settle on a single number of added labeled samples. For this reason, we decide to report the area under the learning curve (ALC) of each active learning method as the criterion. In our particular setting this means, we set an upper bound on the final number of labeled samples to pick and record the accuracy for the ever increasing training set size. Finally, the average over these accuracies is reported as the overall performance measure, which is a criterion, slightly different from the one suggested in (Cawley, 2011), where AUCs are used instead of accuracies and the horizontal axis with the number of added data points is scaled logarithmically.

In addition to the accuracy or its average, we also consider the log-likelihood on the synthetic datasets. The relationship between the performance in terms of the log-likelihood, which is the actual measure that the logistic regressor optimizes, and any secondary performance measure like the error rate, the AUC, or the average accuracy is not necessarily easy to captures. One of the authors has therefore advocated to not only consider accuracies (or any other secondary measure) in an evaluation, but to also consider the loss that the classifier under consideration actually optimizes on the test set. This viewpoint was originally suggested in the semi-supervised learning literature (Loog & Jensen, 2015), but pertains to different learning settings like active learning as well (Loog et al., 2016). Studying the log-likelihood basically gives us a potentially different view on the behavior of the active learners we consider.

3.2. Analysis of synthetic datasets

In the following subsection, we show the results on 3 synthetic data sets, respectively. A preference map is proposed to reveal the characteristics of different active learning algorithms. We also present the effects of class overlapping and feature dimensionality in the end.

3.2.1. Result on Synth1 and Preference Map

Shown in Fig 1, we see the distributions of three synthetic datasets, along with the performance of each active learning method in terms of the classification accuracy and the log-likelihood on the test set. We also present the 90% confidence interval around each learning curve. To start with, note that no single method outperforms all the other methods on all the datasets.

Synth1 is a fairly simple classification problem and some algorithms perform really well in the very beginning stage, such as the variance reduction approaches FIVR and EVR. On the other hand, ENTROPY achieves rather poor performance at the beginning and even the *worst* one at the first selected point in Fig 1b. To understand a specific aspect of how uncertainty sampling acts and generally show a difference in characteristic of the various active learning methods, we introduce what we call the *preference map* (Fig 2). The preference map is obtained by keeping track of the locations of the queried instances selected by each active learning algorithm. Presenting kernel density plots of all these locations and displaying them using pseudo-colors gives an impression where in feature space the active learners request their data from. The highest density regions are marked in red while the lowest density regions are indicated in blue. The preference map of the instance first queried is shown in Fig 2a. We also plot the preference maps corresponding to the complete learning, where we exponentially weigh down later observations based on the intuition that the examples selected early on in the process are more valuable than the examples selected in the

[3] http://qwone.com/ jason/20Newsgroups/

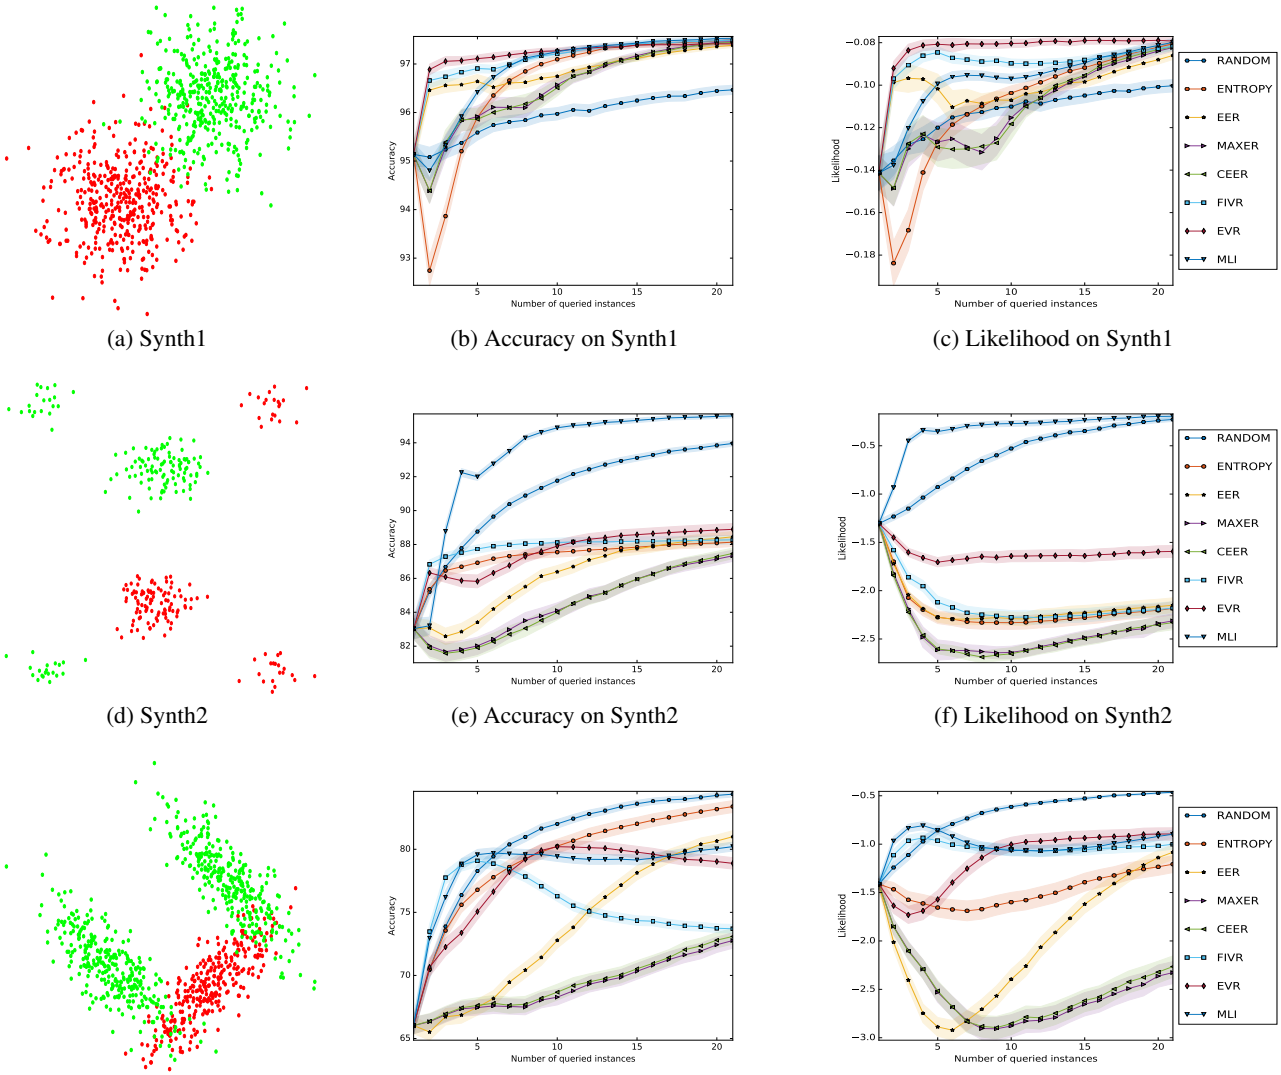

(a) Synth1 (b) Accuracy on Synth1 (c) Likelihood on Synth1

(d) Synth2 (e) Accuracy on Synth2 (f) Likelihood on Synth2

(g) Synth3 (h) Accuracy on Synth3 (i) Likelihood on Synth3

Figure 1. Distribution for three synthetic data sets and the results on these same sets in terms of classification accuracy and log-likelihood with 90% confidence interval. Red and green points represent the two different classes. (a) shows the distribution of the Synth1 dataset; (b) and (c) are the average accuracy and log-likelihood of each active learning method on the test set for Synth1. (d)-(f) and (g)-(i) represent the same results for Synth2 and Synth3, respectively.

Table 2. Data sets information: It shows the number of instances (# INS) and the feature dimensionality (# FEA)

DATA SET	# INS	# FEA	DATA SET	# INS	# FEA
AC-INFLAM	120	6	ACUTE	120	6
AUSTRALIAN	690	14	BLOOD	748	4
BREAST	683	10	CREDIT	690	15
CYLINDER	512	35	DIABETES	768	8
FERTILITY	100	9	GERMAN	1000	24
HABERMAN	306	3	HEART	270	13
HEPATITIS	255	19	HILL	606	100
IONOSPHERE	351	34	LIVER	345	6
MUSHROOMS	1000	112	MAMMOGRAPHIC	961	5
MUSK1	476	166	OOCTRIS2F	912	25
OZONE	1000	72	PARKINSONS	195	22
PIMA	768	8	PLANNING	182	12
SONAR	208	60	SPLICE	1000	60
TICTACTOE	958	9	VC2	310	6
VEHICLE	435	18	WISC	699	9
WDBC	569	31	D VS P	1608	16
E VS F	1543	16	I VS J	1502	16
M VS N	1575	16	V VS Y	1577	16
U VS V	1550	16	3 VS 5	1500	784
5 VS 8	1500	784	7 VS 9	1500	784
BASEBALL VS HOCKEY	1993	500	PC VS MAC	1945	500
MISC VS ATHEISM	1427	500			

later rounds. The specific weight function we employ is $\exp(-2 * r/L)$, where r and L are the current round and the total rounds, respectively. The corresponding preference maps are in Fig 2b.

Looking at Fig 2a, we can see that random sampling of course prefers the region where the mean vector of each class is. Clearly, the preference map for random sampling should ultimately just reproduce the original underlying distribution, which is a mixture of two normal distributions in the setting we consider. Uncertainty sampling clearly prefers to query the points in the middle of two clusters since it focuses on the instances near the estimated decision boundary. Even though these samples may be close to the true decision boundary, they may definitely not be a good choice, because it leads to instable estimates. This is what we see in the results, where ENTROPY performs rather poor in the very beginning stage. CEER and MAXER show similar behavior in the preference map and also seem to give fairly bad performance at the start of the active learning cycle. Their map, however, seems a bit more rectangular, which may lead to slightly improved stability and therefore performance as compared to ENTROPY. It is interesting to see that the variance reduction methods like EVR and FIVR also sample parallel to the decision boundary, but more through the respective class centers, which indeed leads to more stable and therefore better performing estimates. EER displays a bit of the same behavior, but clearly much less than EVR and FIVR. MLI, on the other hand, seems to sample perpendicular to the decision boundary, away from the regions with high density. This may be because MLI, which is similar to QUIER (Huang et al., 2014), tends to balance the informativeness and representativeness. When only two initial points are available, MLI prefers to select the instance far away from already labeled ones. Turning our attention to Fig 2b, we see a dramatic change in behavior for at least four of the strategies. FIVR and especially EVR change their sampling from parallel to more perpendicular to the decision boundary. The changes we see for EER and MLI may be interpreted as a change from the more explorative initial phase to a more exploitative later stage, where a sampling around the decision boundary is performed to refine it. That active learning should actually deal with the exploration-exploitation tradeoff is at the basis of MLI. In addition, we can clearly see from Fig 1b that EER, FIVR, EVR and MLI, of which preference maps vary from and Fig 2a and Fig 2b, significantly outperform the remaining methods.

3.2.2. RESULT ON SYNTH2

Fig 1e shows that on the second synthetic problem, random sampling far surpasses all the active learning methods except for MLI, which is the best performing strategy. Fig 3 illustrates why random sampling can outperform some active learning methods. Assume that we have two initial labeled points which are marked with black color. Each active learning method is allowed to query two more examples. Note that some methods select the same points, such as ENTROPY, MAXER and CEER select the same points which come closest to the decision boundary estimated on the two initial points. However, these selected points can-

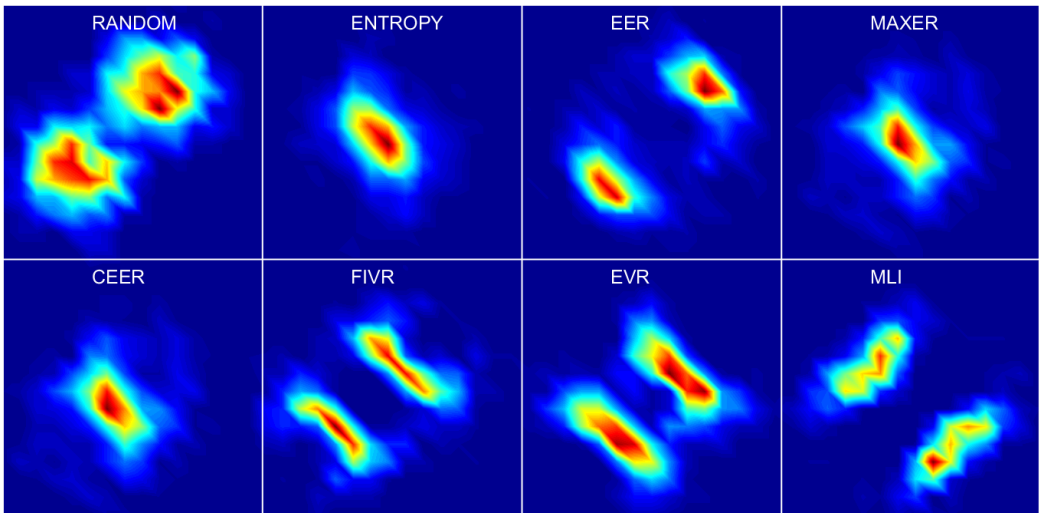

(a)

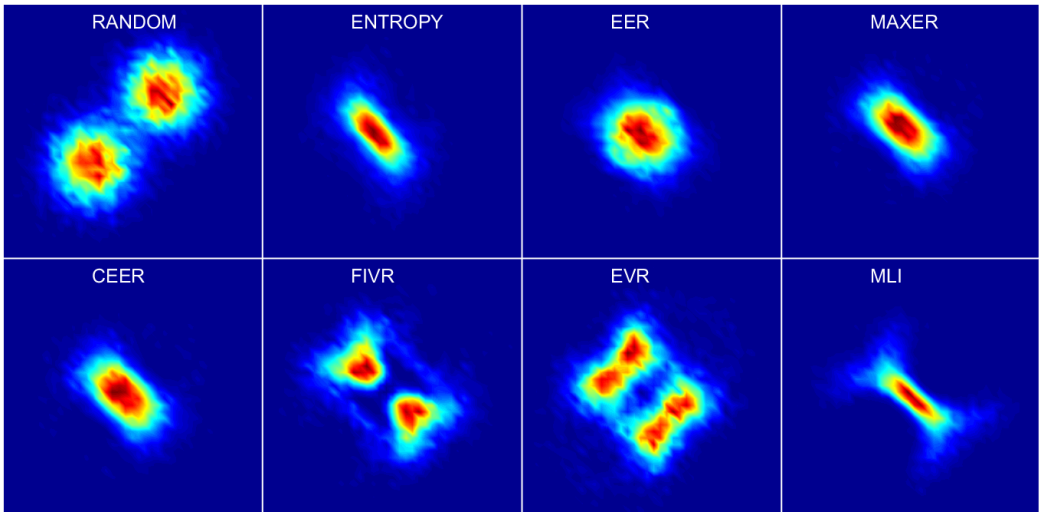

(b)

Figure 2. (a) Preference maps of first queried example selected by each active learning method on the Synth1 dataset; (b) Weighted preference maps over the whole learning process on the same problem.

not provide much more information about the true distribution. Take uncertainty sampling for example, it keeps querying the points closest to the initial decision boundary estimated from the limited labeled instances. In this setting, the two initial points lead to a completely wrong estimation of the decision boundary. This shows a common situation where some active learning methods get stuck in keeping querying useless instances due to inaccurate estimation of model parameter. Random sampling does not suffer from this because it acts purely random in selecting new instances.

MLI can perform even better in this situation (Huang et al., 2014). MLI selects two examples which are far away from the wrongly estimated decision boundary since it also considers the so-called representativeness of each instance, such as whether the instances are inside of some clusters. This leads to the exploratory behavior of MLI. To demonstrate the clearly different behavior of MLI, Fig 4 displays preference maps corresponding to the whole learning curve on Synth2 dataset. MLI is more likely to query the instances along four clusters on the border line, while some methods like uncertainty sampling and error reduction approaches favor the instances in the two middle clusters. This is the reason why MLI can significantly outperform random sampling and other active learning methods on this artificial set.

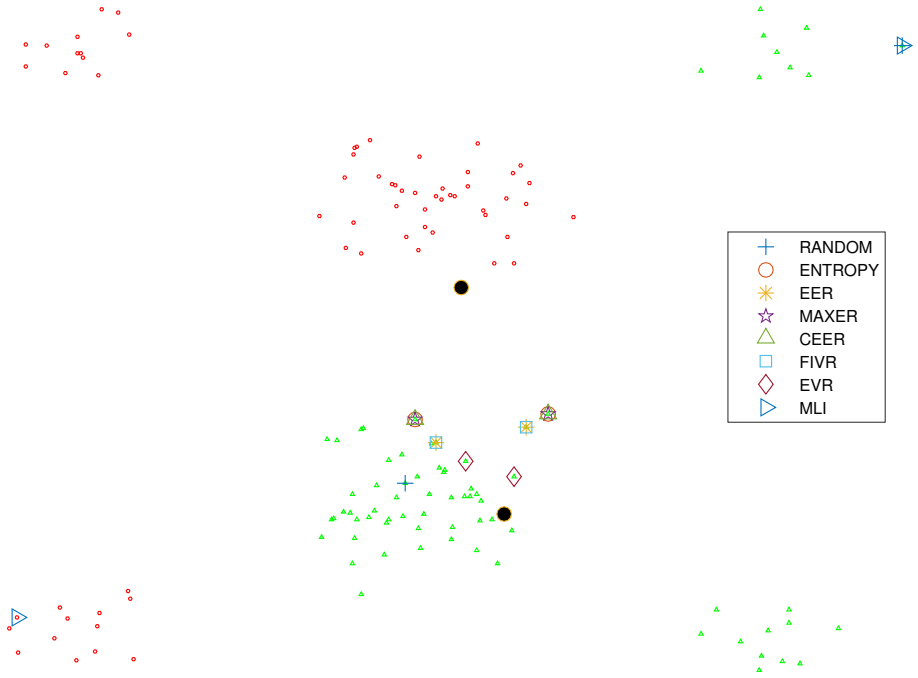

Figure 3. Queried instances by different active learning methods. Two black points represent two initial labeled points. Two more instances are queried by each active learning method.

3.2.3. Result on Synth3

From Fig 1h, we can observer some negative result that random sampling outperforms all the other active learning methods after 5 instances are selected. The possible reason is that random sampling can explore the whole structure of this dataset while other methods just pay attention to some certain parts without exploring the whole dataset. And another reason may be that this dataset is a little difficult due to its complex structure. On this kind of hard datasets, active learning methods can easily get stuck in local structure while ignoring the global view of the problem. Due to space limitations, the preference maps of Synth3 are omitted. From Fig 1, we also find that MAXER and CEER always get a poor performance compared with EER, EVR and MLI, which indicates the best-case criteria is not reasonable in active learning.

As indicated already, we also compute the log-likelihood the varies strategies attain on the test set. Looking at Fig 1c and Fig 1b, the first main observation we can make is that a better likelihood—and therefore, in the sense of logistic regression, a better model—does definitely not necessarily lead to improved classification performance. A second important observation is that in terms of the log-likelihood, many active learners achieve a strongly decreased likelihood, while the accuracy may increase. In some settings, it is the other we around and we see an increasing likelihood, while the accuracy drops (e.g., see the cruves related to EVR on Synth3).

3.2.4. Class overlapping and feature dimensionality

We also explore how the class overlapping and feature dimensionality affect the performance of each active learning algorithm. As is shown in Fig 5, we show the results on four different overlap rates. Four colours represent four different overlap rates, and the numerical value behind each colour indicts the Bayes error of the classification problem. The value in vertical axis represents the area under the curve of the whole active learning procedure. Red dashed line shows the best achieved performance among all the methods; black dashed line corresponds to the performance of random sampling. It is clear that the higher the overlapping, the lower the scores of all active learning methods achieve. It is reasonable since higher overlapping indicates more difficult classification problem. We can also find that FIVR behaves robust in all overlapping setting, especially in the higher overlap rate situation, which indicts that FIVR may work well on difficult tasks. Error reduction approaches are even outperformed by random sampling in the higher overlapping setting, as well as uncertainty sampling. We also adjust the dimensionality of Synth1 dataset from 2 to 20 and 200 dimensions to investigate the influence of feature dimensionality. The results are omitted due to the space limit, however, similar results are observed in the experiment. We find that the higher the dimensionality, the

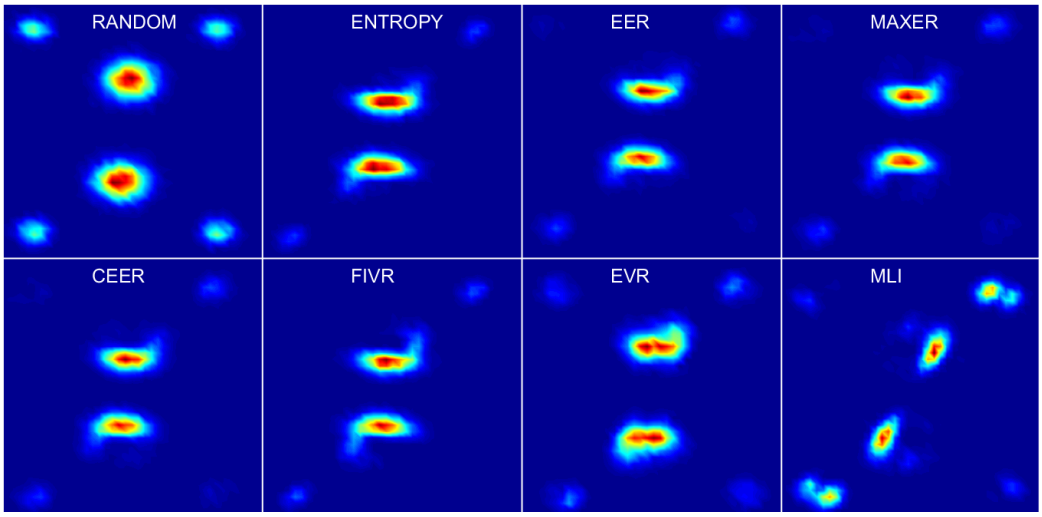

Figure 4. Preference maps for the whole learning procedure by each of the active learning methods on the Synth2 dataset.

worse the performance of each method. One cause may be that high dimensionality problems need more training samples to achieve the same performance. It also demonstrates that ENTROPY behaves worse along with the increase of feature dimensionality. For example, when the feature is 200-dimensional, ENTROPY behaves worse than random sampling.

3.3. Analysis on real-world datasets

Table 3 presents the results for each active learning method on the 43 real-world datasets. We adopt the paired t-test at a 95% significance level on all the experiments to test which method does not significantly differ from the best method. The best performance is highlighted in bold face, together with the competitors that perform at a comparable level. The average accuracy ("Mean" in Table 3) of each method is also reported. "Win counts" shows the total number of datasets on which one method achieves the best performances. "win/tie/loss vs. random" demonstrates the win, tie, and loss counts of one method versus random sampling over all of the datasets.

As is shown in Table 3, no algorithms outperform all others on all the datasets. Still ENTROPY, EER and MLI seem to markedly outperform other active learning methods. It may be surprising that uncertainty sampling can compete with relatively sophisticated active learning algorithms as it is a rather simplistic approach. In fact, uncertainty sampling gets the best accuracy and performs best in "win/tie/loss" counts versus random sampling. MLI behaves the best among the remaining methods in terms of "win counts" while its average accuracy is outperformed by uncertainty sampling and EER.

Considering the error reduction approaches, it is clear that EER outperforms MAXER and CEER. The overall performances of MAXER and CEER remain close to that of random sampling. MAXER merely surpasses random sampling on 19 of the 43 datasets. This seems to demonstrate that the best-case criterion is not a reasonable choice for active learning, at least for error reduction approaches. The possible cause may be that such optimistic measure simply puts too much trust in a typically badly estimated model. As a result, initial errors may get reinforced rather than mitigated by correctly chosen additional samples. This is comparable to some of the issues that arise in self-learning and EM-based approaches to semi-supervised learning (Loog & Jensen, 2015; Loog et al., 2016; Cohen et al., 2004). (Guo & Greiner, 2007) proposed the on-line adjustment for MAXER, which switches to another active learning method when MAXER "guesses" wrong about the true label of latest queried instance. We do not adopt this adjustment since it can be used for any active learning algorithms and we care about the performance of pure active learning methods. CEER obtains performance comparable to MAXER, and it shares the same problem with MAXER since it also use the optimistic strategy (Guo & Schuurmans, 2008). One possible reason why CEER can not work is that the trade-off parameter α is not well determined. Compared with error reduction approaches, MLI performs worse than EER while it is far better than MAXER and CEER. As to the variance reduction approaches, EVR seems to slightly outperform FIVR in terms of average scores and "win/tie/loss" counts. FIVR which achieves better performance than MAXER and CEER, however, is still exceed by random sampling on 12 out of 43 datasets. EVR behaves comparable with EER.

As it is the technique to beat, it is important to see how the active learners perform in comparison with random sam-

Table 3. Performance Comparison on the Area under the Learning Curve (Accuracy). The higher the score, the better the performance. For each data set, the best performances and its comparable competitors are highlighted in bold face. Average performance of all the active learning methods are also reported as "mean". "Win counts" shows the total numbers of datasets on which one method achieves the best performances. "win/tie/loss vs. random" demonstrates the win/tie/loss counts of one method versus random sampling on the whole datasets based on paired t-test at 95 percent significance level.

DATASET	RANDOM	ENTROPY	EER	MAXER	CEER	FIVR	EVR	MLI
HILL	0.581	0.593	0.616	0.606	0.575	0.549	0.59	**0.626**
PLANNING	0.586	0.584	0.58	0.568	0.574	0.574	0.587	**0.614**
CYLINDER	0.586	0.593	0.61	0.575	0.576	**0.63**	0.591	0.608
LIVER	0.627	0.612	0.635	0.623	0.621	**0.645**	0.632	0.615
SPLICE	0.659	**0.685**	0.679	0.666	0.663	0.676	0.664	0.65
GERMAN	0.664	**0.694**	0.673	0.652	0.654	0.687	0.678	0.691
OOCTRIS2F	0.679	0.665	0.678	0.648	0.652	0.651	0.68	**0.686**
MUSK1	0.682	**0.702**	**0.699**	0.672	0.668	0.679	0.689	**0.702**
FERTILITY	0.693	0.701	0.706	0.679	0.674	0.716	0.686	**0.727**
HABERMAN	**0.711**	**0.712**	**0.712**	0.704	0.704	0.692	0.708	0.694
SONAR	0.713	0.711	0.715	0.713	0.715	**0.72**	**0.72**	0.708
PIMA	0.716	0.717	0.706	0.71	0.707	**0.727**	0.708	0.711
PCMAC	0.717	0.727	0.715	0.698	0.696	0.717	0.711	**0.747**
DIABETES	0.719	0.721	0.723	**0.726**	0.724	0.707	**0.725**	**0.726**
RELIGIONATHEISM	0.72	**0.74**	0.71	0.687	0.687	0.712	0.704	0.691
HEPATITIS	0.731	**0.753**	**0.753**	0.744	0.741	0.745	**0.754**	0.73
BLOOD	**0.743**	0.718	0.74	0.732	0.736	0.723	0.732	0.73
HEART	0.774	0.793	0.791	0.788	0.788	**0.799**	0.781	0.797
IONOSPHERE	0.779	0.782	**0.818**	0.806	0.801	0.79	0.812	0.674
CREDIT	0.779	**0.822**	0.793	0.795	0.804	**0.819**	0.791	0.797
MAMMOGRAPHIC	0.78	0.779	0.774	0.781	0.784	**0.795**	0.777	0.766
BASEHOCKEY	0.793	**0.822**	0.784	0.772	0.77	**0.82**	0.783	0.817
VC2	0.807	0.814	0.815	0.802	0.803	0.816	0.822	**0.825**
PARKINSONS	0.811	0.823	0.824	0.824	**0.828**	0.825	0.825	**0.83**
AUSTRALIAN	0.823	**0.844**	0.832	0.839	0.838	0.817	0.831	0.842
LETTERIJ	0.853	0.871	**0.879**	0.807	0.806	0.841	0.869	0.865
LETTERVY	0.855	0.88	0.878	0.814	0.814	0.753	**0.886**	0.861
3VS5	0.856	0.884	**0.903**	0.89	0.889	0.869	**0.903**	0.859
VEHICLE	0.859	**0.884**	0.878	0.851	0.855	0.83	**0.884**	**0.883**
5VS8	0.864	0.891	**0.907**	0.896	0.895	0.875	**0.909**	0.85
7VS9	0.876	0.904	0.914	0.905	0.906	0.909	**0.917**	0.841
OZONE	0.882	0.884	0.86	0.862	0.861	0.843	**0.901**	0.892
TICTACTOE	0.894	0.902	**0.912**	0.673	0.684	0.903	0.898	0.853
LETTERMN	0.916	**0.939**	**0.944**	0.91	0.91	0.932	0.941	0.927
MUSHROOMS	0.931	**0.971**	0.969	0.967	0.967	**0.972**	0.96	**0.971**
LETTEREF	0.933	**0.959**	0.954	0.949	0.95	0.954	0.956	0.956
WDBC	0.938	0.955	0.953	0.955	0.954	0.943	0.951	**0.958**
LETTERDP	0.939	0.964	0.963	0.95	0.95	0.954	0.961	**0.967**
LETTERUV	0.945	0.97	**0.972**	0.955	0.955	0.963	0.966	**0.974**
WISC	0.949	0.956	0.951	**0.958**	0.957	0.954	0.953	0.956
BREAST	0.95	0.958	0.956	0.956	0.957	0.96	0.955	**0.962**
AC-INFLAM	0.955	**0.985**	0.981	0.962	0.965	0.982	0.967	0.98
ACUTE	0.978	**0.991**	0.971	0.96	0.966	**0.991**	0.956	**0.992**
MEAN	0.796	**0.811**	**0.810**	0.791	0.791	0.801	0.807	0.804
WIN COUNTS	2	15	10	2	1	10	9	**16**
WIN/TIE/LOSS VS. RANDOM	-	**32/7/4**	31/3/9	19/5/19	20/3/20	29/2/12	31/3/9	29/2/12

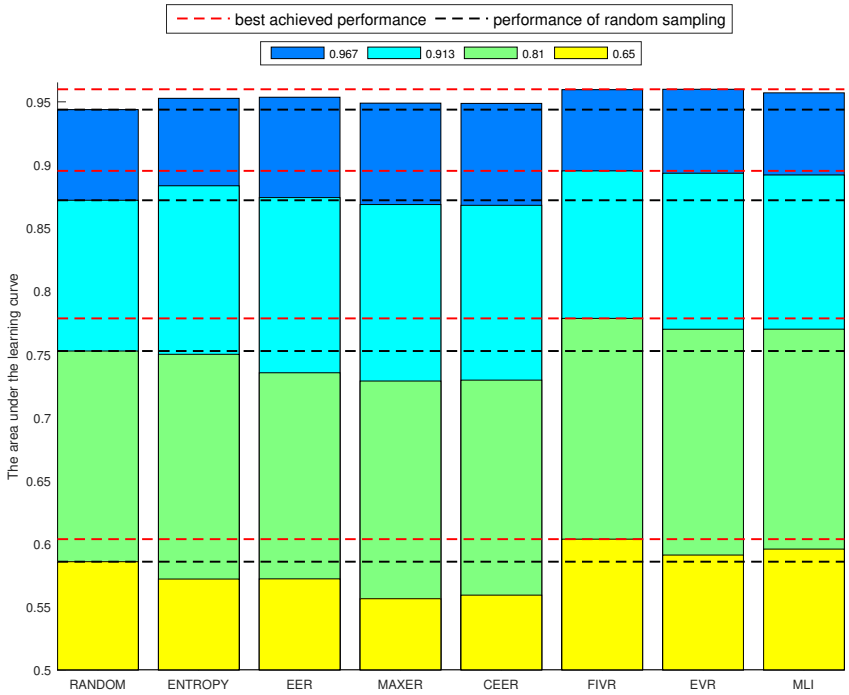

Figure 5. Comparisons on different overlapping rates. Different color represent four overlapping setting, following with numerical values which indict the Bayes error. Red dashed line describes the best achieved performance among all the methods; black dashed line corresponds to the performance of random sampling.

pling over the 43 datasets. We therefore consider the ratio $\frac{V_{active}}{V_{random}}$ with V_{active} and V_{random} the performance of active learning and random sampling, respectively. This gives us an indication of the relative improvement (or deterioration) the active learning schemes provide. We compute the ratios over all the datasets and visualize the outcomes with a box plot in Fig 6. The 25th, 50th and 75th percentiles are shown and the magenta crosses indict the average values of the ratios. We can observe that ENTROPY and EER may deliver satisfactory performances, while MAXER and CEER behave rather poor. MLI achieves a highest ratio on one dataset, which means that MLI can improve most upon random sampling sometimes. Possibly more important, however, ENTROPY and EER may be considered safer: they may not reach the relative improvements that MLI achieves, but at least they also do not display dramatic decreases in performance.

Table 3 is divided into three different sections according to roughly indicate the difficulty level of each dataset, i.e., how high an accuracy random sampling achieves. The first group, which arranges from 0.5 to 0.78, represents the datasets on which reaching good accuracies seems difficult. The second group, arranging from 0.78 to 0.90, corresponds to the datasets which have medium accuracies apply. The last group consists of the remaining datasets, which seem fairly easy to solve by a linear logistic classifier. We can see that random sampling surpasses all the other methods on blood data sets, which only occurs in the difficult classification tasks. For the medium and easy datasets, random sampling does not achieve the best performances any more, which may indicate that we need only consider random sampling on relatively hard tasks. For the difficult classification problems in the first group, ENTROPY, FIVR and MLI achieve comparable performances. FIVR is more likely to work on the hard tasks than on the easy tasks. This matches with the results in Fig 5 for the experiments where the ratio of overlap is varied. When the overlap ratio increases, the difficulty-level gains, and FIVR improves over other methods. In the second group, EVR obtains the best performance, while it seems to underperform in the other groups. For the easy datasets, it seems that MLI and ENTROPY get slightly better than the other methods. ENTROPY also performs well on the hard and medium datasets. The experiments demonstrate that uncertainty sampling is a robust active learning algorithm, regardless of the difficulty-level of the tasks. EER is a runner-up which also performs robust in the overall comparison.

3.4. Computational Cost Analysis

Computational cost can also be a critical issue when employing active learning methods. Table 4 assesses the average computational cost of each of the methods. All the

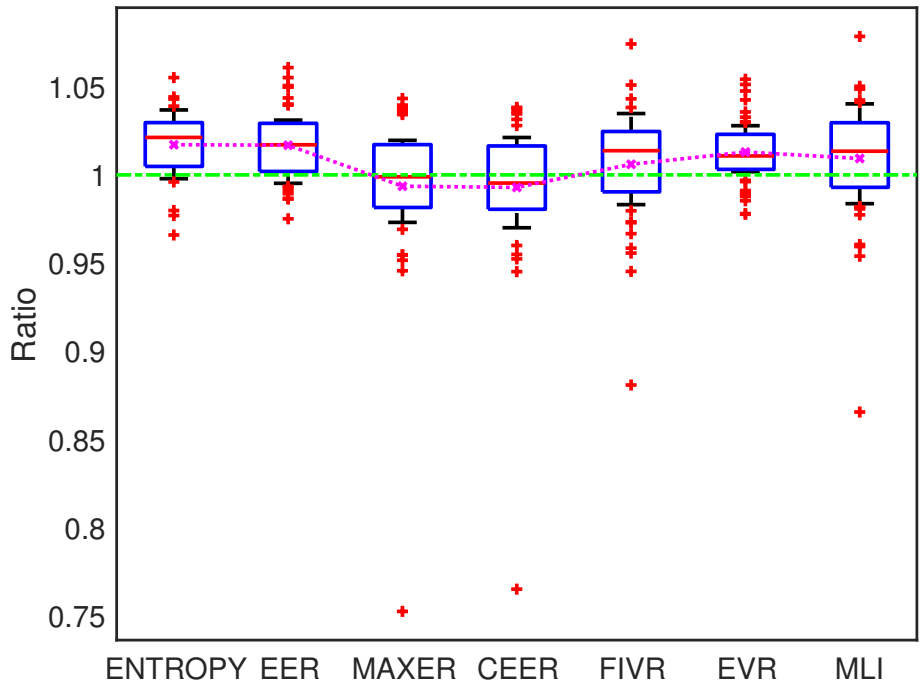

Figure 6. Box plots of the ratios of active learning to random sampling over all the datasets. The magenta crosses represent the average values of the ratios. The green line is at one, at the performance of random sampling.

Table 4. Computational cost comparison of each query for each active learning method (in seconds)

DATASET (#INS #FEA)	RANDOM	ENTROPY	EER	MAXER	CEER	FIVR	EVR	MLI
ACUTE (120,6)	0.0002	0.0004	0.0217	0.0218	0.0319	0.0028	0.0285	0.0197
AUSTRALIAN (690,14)	0.0001	0.0005	0.2247	0.2247	0.2791	0.0221	0.4148	0.1308
MUSK1 (476, 166)	0.0001	0.0024	0.8845	0.8852	0.9190	0.3087	1.5363	0.1520
HILL (600, 100)	0.0001	0.0020	0.9138	0.9474	1.0111	0.1096	1.3417	0.1736
MUSHROOMS (1000, 112)	0.0001	0.0009	0.5547	0.5367	0.5954	0.1848	2.2543	0.2363
LETTEREF (1543, 16)	0.0001	0.0009	0.8824	0.9231	1.0670	0.0461	1.7103	0.3296
3VS5 (1500, 784)	0.0002	0.0090	10.5462	10.5235	9.1991	45.3676	131.7279	1.7962
BASEHOCKEY (1993, 500)	0.0001	0.0265	64.2465	65.0460	67.9577	19.1748	128.9946	2.5100
MEAN	0.0001	0.0053	9.7843	9.8886	10.1325	8.1521	33.5010	0.6685

experiments are constructed with MATLAB 8.6 on a Intel(R) Core(TM) i7-4600U CPU PC with 8 GB memory. We test the computational cost on 8 datasets that vary in the numbers of instances and the feature dimensionalities. Clearly, random sampling and ENTROPY are the most efficient methods due to their simplicity. MLI also has a low computational burden compared with other algorithms. Error reduction and variance reduction have, on the other hand, a significantly higher computational cost than other methods. Both of them are especially significantly less efficient on the high dimensionality datasets like 3vs5 and basehockey. The reason may be that both need to retrain the logistic regressor in every selection step over all the unlabeled instances and all possible labels, which is relatively time consuming especially in higher dimensions. We also note that EVR has the highest computational cost. This is because EVR has to repeatedly calculate the inverse of matrix, which is extremely time consuming.

4. Conclusions

We compared the objective functions employed and the performance level attained for current state-of-the-art active learning methods for logistic regression. We also pointed out the main relationships between these methods. The experiments on the synthesis datasets and the large number of real-world datasets show some of the chief underlying characteristics of each of the active learning methods. On the average, we would deem ENTROPY and EER the more promising methods. Though ENTROPY is a rather simplistic criterion and quite short-sighted when picking instances, it does outperform min-max view ap-

proach and variance reduction methods in our experiments. MLI demonstrates its advantage in querying the representative instances on the synthesis data Synth2. We also observe that FIVR is one hopeful active learning method especially on complex and difficult tasks. A possible downside of expected error reduction approaches is the high computational cost it incurs. Variance reduction approaches and MLI also suffer the same problem. How to speed up these methods may be one problem for further research.

Overall, on the positive side, we can conclude that active learning can indeed provide improved performance over random sampling, most certainly if we consider the whole ensemble of active learners in this work. This, however, also seems to be the big negative. On its own, none of these methods can prevent becoming worse than random sampling. While this seems impossible anyway for every single instantiation of a problem, our results indicate that it does not even hold in the average. That is, for every method there seem to be (real-world) datasets for which the method performs significantly worse than random sampling, even when averaged over multiple runs. Finding active learning methods that are, in some sense, safe and yet give significant performance gains at times, seems to be the challenge ahead.

References

Atkinson, Anthony, Donev, Alexander, and Tobias, Randall. *Optimum experimental designs, with SAS*. Oxford Univ. Press, UK, 2007.

Cai, Deng and He, Xiaofei. Manifold adaptive experimental design for text categorization. *Knowledge and Data Engineering, IEEE Transactions on*, 24(4):707–719, 2012.

Cai, Wenbin, Zhang, Ya, Zhou, Siyuan, Wang, Wenquan, Ding, Chris, and Gu, Xiao. Active learning for support vector machines with maximum model change. In *Machine Learning and Knowledge Discovery in Databases*, pp. 211–226. Springer, 2014.

Cawley, Gavin C. Baseline methods for active learning. In *Active Learning and Experimental Design@ AISTATS*, pp. 47–57, 2011.

Chakraborty, Shayok, Balasubramanian, Vineeth, and Panchanathan, Sethuraman. An optimization based framework for dynamic batch mode active learning. *Advances in Neural Information Processing*, 2010.

Chakraborty, Shayok, Balasubramanian, Vineeth, and Panchanathan, Sethuraman. Adaptive batch mode active learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2014.

Chang, Chih-Chung and Lin, Chih-Jen. Libsvm: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):27, 2011.

Cohen, Ira, Cozman, Fabio G, Sebe, Nicu, Cirelo, Marcelo C, and Huang, Thomas S. Semisupervised learning of classifiers: Theory, algorithms, and their application to human-computer interaction. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 26 (12):1553–1566, 2004.

Cuong, Nguyen Viet, Lee, Wee Sun, and Ye, Nan. Near-optimal adaptive pool-based active learning with general loss. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2014.

Fan, Rong-En, Chang, Kai-Wei, Hsieh, Cho-Jui, Wang, Xiang-Rui, and Lin, Chih-Jen. Liblinear: A library for large linear classification. *The Journal of Machine Learning Research*, 9:1871–1874, 2008.

Fernández-Delgado, Manuel, Cernadas, Eva, Barro, Senén, and Amorim, Dinani. Do we need hundreds of classifiers to solve real world classification problems? *The Journal of Machine Learning Research*, 15(1):3133–3181, 2014.

Frey, Peter W and Slate, David J. Letter recognition using holland-style adaptive classifiers. *Machine Learning*, 6: 161, 1991.

Freytag, Alexander, Rodner, Erik, and Denzler, Joachim. Selecting influential examples: Active learning with expected model output changes. In *Computer Vision–ECCV 2014*, pp. 562–577. Springer, 2014.

Grandvalet, Yves and Bengio, Yoshua. Semi-supervised learning by entropy minimization. In *Advances in neural information processing systems*, pp. 529–536, 2004.

Gu, Quanquan, Zhang, Tong, and Han, Jiawei. Batch-mode active learning via error bound minimization. *Urbana*, 51:61801, 2014.

Guo, Yuhong. Active instance sampling via matrix partition. In *Advances in Neural Information Processing Systems*, pp. 802–810, 2010.

Guo, Yuhong and Greiner, Russell. Optimistic active-learning using mutual information. In *IJCAI*, volume 7, pp. 823–829, 2007.

Guo, Yuhong and Schuurmans, Dale. Discriminative batch mode active learning. In *Advances in neural information processing systems*, pp. 593–600, 2008.

Guyon, Isabelle, Cawley, Gavin C, Dror, Gideon, and Lemaire, Vincent. Results of the active learning challenge. *Active Learning and Experimental Design@ AISTATS*, 16:19–45, 2011.

Hoi, Steven CH, Jin, Rong, Zhu, Jianke, and Lyu, Michael R. Batch mode active learning and its application to medical image classification. In *Proceedings of the 23rd international conference on Machine learning*, pp. 417–424. ACM, 2006.

Hoi, Steven CH, Jin, Rong, Zhu, Jianke, and Lyu, Michael R. Semi-supervised svm batch mode active learning for image retrieval. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pp. 1–7. IEEE, 2008.

Holub, Alex, Perona, Pietro, and Burl, Michael C. Entropy-based active learning for object recognition. In *Computer Vision and Pattern Recognition Workshops, 2008. CVPRW'08. IEEE Computer Society Conference on*, pp. 1–8. IEEE, 2008.

Huang, Sheng-Jun, Jin, Rong, and Zhou, Zhi-Hua. Active learning by querying informative and representative examples. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(36):1936–1949, 2014.

Kanamori, Takafumi. Pool-based active learning with optimal sampling distribution and its information geometrical interpretation. *Neurocomputing*, 71(1):353–362, 2007.

Kanamori, Takafumi and Shimodaira, Hidetoshi. Active learning algorithm using the maximum weighted log-likelihood estimator. *Journal of Statistical Planning and Inference*, 116(1):149–162, 2003.

Lang, Ken. Newsweeder: Learning to filter netnews. In *Proceedings of the Twelfth International Conference on Machine Learning*, pp. 331–339, 1995.

LeCun, Yann, Bottou, Léon, Bengio, Yoshua, and Haffner, Patrick. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Lewis, David D and Gale, William A. A sequential algorithm for training text classifiers. In *Proceedings of the 17th annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 3–12. Springer-Verlag New York, Inc., 1994.

Li, Chun-Liang, Ferng, Chun-Sung, and Lin, Hsuan-Tien. Active learning with hinted support vector machine. In *ACML*, pp. 221–235, 2012.

Lichman, M. UCI machine learning repository, 2013. URL http://archive.ics.uci.edu/ml.

Liu, Anqi, Reyzin, Lev, and Ziebart, Brian D. Shift-pessimistic active learning using robust bias-aware prediction. In *AAAI*, pp. 2764–2770, 2015.

Loog, Marco and Jensen, Are Charles. Semi-supervised nearest mean classification through a constrained log-likelihood. *Neural Networks and Learning Systems, IEEE Transactions on*, 26(5):995–1006, 2015.

Loog, Marco, Krijthe, Jesse Hendrik, and Jensen, Are Charles. On measuring and quantifying performance: Error rates, surrogate loss, and an example in ssl. In Chen, C. H. (ed.), *Handbook of Pattern Recognition and Computer Vision*, chapter 1.3, pp. 53–68. World Scientific, 5 edition, 2016.

McCallumzy, Andrew Kachites and Nigamy, Kamal. Employing em and pool-based active learning for text classification. In *Proc. International Conference on Machine Learning (ICML)*, pp. 359–367. Citeseer, 1998.

Roy, Nicholas and McCallum, Andrew. Toward optimal active learning through monte carlo estimation of error reduction. *ICML, Williamstown*, 2001.

Schein, Andrew I and Ungar, Lyle H. Active learning for logistic regression: an evaluation. *Machine Learning*, 68 (3):235–265, 2007.

Settles, Burr. Active learning literature survey. *University of Wisconsin, Madison*, 52(55-66):11, 2010.

Settles, Burr, Craven, Mark, and Ray, Soumya. Multiple-instance active learning. In *Advances in neural information processing systems*, pp. 1289–1296, 2008.

Seung, H Sebastian, Opper, Manfred, and Sompolinsky, Haim. Query by committee. In *Proceedings of the fifth annual workshop on Computational learning theory*, pp. 287–294. ACM, 1992.

Shannon, Claude Elwood. A mathematical theory of communication. *Bell System Technocal Journal*, 27(1):379–423, 623–656, 1948.

Tang, Min, Luo, Xiaoqiang, and Roukos, Salim. Active learning for statistical natural language parsing. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, pp. 120–127. Association for Computational Linguistics, 2002.

Tong, Simon and Koller, Daphne. Support vector machine active learning with applications to text classification. *The Journal of Machine Learning Research*, 2:45–66, 2002.

Yu, Kai, Bi, Jinbo, and Tresp, Volker. Active learning via transductive experimental design. In *Proceedings of the 23rd international conference on Machine learning*, pp. 1081–1088. ACM, 2006.

Zhang, Tong and Oles, F. The value of unlabeled data for classification problems. In *Proceedings of the Seventeenth International Conference on Machine Learning,(Langley, P., ed.)*, pp. 1191–1198. Citeseer, 2000.

Zhu, Jingbo, Wang, Huizhen, Tsou, Benjamin K, and Ma, Matthew. Active learning with sampling by uncertainty and density for data annotations. *IEEE Transactions on Audio, Speech, and Language Processing*, 18(6):1323–1331, 2010.

Zhu, Xiaojin, Lafferty, John, and Ghahramani, Zoubin. Combining active learning and semi-supervised learning using gaussian fields and harmonic functions. In *ICML 2003 workshop on the continuum from labeled to unlabeled data in machine learning and data mining*, pp. 58–65, 2003.